

APPLICATION  
FOR  
UNITED STATES LETTERS PATENT

TITLE: AUTOMATIC POWER DOWN

APPLICANT: DAVID C. CHALUPSKY, GREGORY V. GRITTON AND  
THOMAS L. STACHURA

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL485674196US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit February 22, 2002

Signature 

Leroy Jenkins  
Typed or Printed Name of Person Signing Certificate

## Automatic Power Down

### **TECHNICAL FIELD**

This invention relates to power management of network devices.

### **BACKGROUND**

5           Industry establish hardware, software and data structure interfaces which, when implemented, enable an operating system to manage power consumption of peripheral devices and buses. Communication devices such as Ethernet network devices often incorporate industry standards giving the devices the ability to  
10           be powered down to a sleep state or an off state by the operating system of the computer.

### **DESCRIPTION OF DRAWINGS**

FIG. 1 is a diagram of a network device connected to a host computer system over a PCI Bus.

15           FIG. 2 is a diagram of a Gigabit Ethernet device with an integrated physical layer interface.

FIG. 3 is a diagram of power management state transitions of the Gigabit Ethernet device of FIG. 2.

**DETAILED DESCRIPTION**

Referring to FIG. 1, a communication arrangement 10 includes a host computer system 12 and a network device 16, e.g., a Gigabit Ethernet device, that both interface with a Peripheral Component Interconnect (PCI) bus 14. The host computer system 12 executes an operating system 12a stored on a computer readable medium (not shown) and loaded into resident memory (not shown) when the host computer system 12 boots up. The operating system 12a controls functions of the host computer system 12 and could cause the host computer system 12 or the network device 16 to enter a sleep or low power state. The host computer system includes a main power supply 12b and an auxiliary power supply 12c. The network device 16 includes a physical layer interface 18 that interfaces with a physical network link 20.

Referring to FIG. 2, a Gigabit Ethernet device 16 interfaces with a PCI bus 14 and a physical network link 20. The Gigabit Ethernet device 16 includes a Media Access Control (MAC) subsystem 30 and an electrically powered physical layer interface 18. The MAC subsystem 30 includes a PCI Bus Interface 32, a Direct Memory Access (DMA) Controller 36, and a Media Access Controller (MAC) 38. The PCI Bus Interface 32 includes a 16-bit Power Management Control/Status (PMCS) Register 34 which includes a two-bit Power State field (not shown) used both to

determine the current power state of the device 16 by the operating system 12a and to permit the operating system 12a to set the device 16 to a different power management state in accordance with the definitions provided in Table I.

5

Table I

| <b>Bit Value</b> | <b>Power Management State of Device 16</b> | <b>Power Management State Definition (consistent with Advanced Configuration and Power Interface (ACPI) Specification, Revision 2.0, published July 27, 2000)</b> |
|------------------|--|---|
| 00               | D0   | Device 16 is on and running. It is receiving full power from the host computer system 12 and is delivering full functionality to the user.                        |
| 11               | D3   | Device 16 is off.   |

Other embodiments may support the additional ACPI Specification 2.0 power management states D1 and D2, in which case the operating system may write the additional bit values set out in Table II to transition the device 16 to power management state D1 or D2.

10

Table II

| <b>Bit Value</b> | <b>Power Management State</b> | <b>Power Management State Definition (consistent with ACPI Specification 2.0)</b>   |
|------------------|-------------------------------|---|
| 01               | D1                            | Device 16 is in a low-power sleep state which device context may or may not be lost.  |
| 10               | D2                            | Device is in a low-power state at attains greater power savings than state D1 in which device context may or may not be lost. |

Referring to FIG. 3, a power management state transition 50 of the Gigabit Ethernet device 16 depicted in FIG. 2 is shown. The Gigabit Ethernet device 16 has four power management states, D0 Active (52), D0 Uninitialized(54), D Reset (56) and D3 (58) and they are generally defined in Table III.

Table III

| <b>Power Management State of Device 16</b> | <b>Definition</b>  |
|--|--|
| D0 Active                                  | Device 16 is on and running and is delivering full functionality and performance to the user.              |
| D0 Uninitialized                           | Device 16 is powering up and awaiting initialization by Basic Input/Output System (BIOS).                  |
| D Reset                                    | A PCI reset signal is asserted and the device 16 is and is awaiting a deassertion of the PCI reset signal. |
| D3   | Device 16 is off.  |

D0 Active (52) and D0 Uninitilized (54) are subsets of the ACPI Specification 2.0 power management state D0. Additionally, the Gigabit Ethernet device's 16 power management state D3 (58) includes the D3<sub>hot</sub> state, which means that the Gigabit Ethernet device 16 can be transitioned to the D0 Uninitialized (54) state via software by writing "00" to the device's 16 Power Management Control/Status register 34 or by having the PCI reset signal asserted. The Gigabit Ethernet device 16 power management state

D3 (58) also includes the D3<sub>cold</sub> power management state, which means that the Gigabit Ethernet device 16 is transitioned to a D0 Uninitialized state (54) by reapplying the main power supply, 12b and deasserting the PCI reset signal. However, other  
5 embodiments may only support the D3<sub>cold</sub> power management state or the D3<sub>hot</sub> power management state.

The Gigabit Ethernet device 16 may transition to a different power management state as a result of the operating system 12a or Basic Input/Output System (BIOS) directing a power management state change (60, 62, 64). For example, if the  
10 Gigabit Ethernet device 16 is up and running in the D0 Active state (52) and the operating system 12a writes "11" to the Power State field in the Power Management Control/Status Register (62), the device 16 transitions to state D3 (58). Similarly, if  
15 the device 16 is in power management state D3 (58) and the operating system 12a writes "00" to the Power State field in the Power Management Control/Status Register (62), the device 16 transitions to state D0 Uninitialized (54). The device 16 also transitions from D0 Uninitialized (54) to the D0 Active (52)  
20 when the BIOS writes a "1" to the memory access enable bit (64) of a PCI command register (not shown) on the PCI interface 32.

In addition to changing power management states as a result of direction from the operating system (60, 62) or BIOS (64), the Gigabit Ethernet device 16 also automatically changes power

management states upon detecting an assertion or deassertion of the PCI reset signal (66, 68) or an assertion of the chip reset signal (71).

An assertion of a PCI reset signal (PCI\_RST#) occurs  
5 shortly before the host computer system 12 transitions from its main power supply 12b to its auxiliary power supply 12c.

The PCI bus interface 32 is configured to monitor the PCI reset signal on the PCI bus 14. If an assertion of the PCI reset signal (PCI\_RST#) is detected (66), the MAC subsystem 30  
10 automatically reverts the device 16 to a reset power state, D Reset (56). The device 16 remains in state D Reset (56) until the PCI interface 32 senses a deassertion of the PCI reset signal (68), which causes the MAC subsystem 30 to transition the device 16 to the D0 Uninitialized state (54).

15 The device 16 is also in the D Reset state (56) when it receives a chip reset signal (LAN\_PWR\_GOOD assertion) (71) the first time the host computer system 12 receives power after power down of both the main 12b and auxiliary 12c power supplies.

20 Referring to FIGS. 1-3, the MAC subsystem 30 controls the power to the physical layer interface 18 dependent upon the power management state of the device 16 and whether wake-up is enabled. A one-bit Power Management Event Enable (PME\_En) field within the PMCS Register 34 indicates whether wake up has been

enabled. Wake up may be enabled or disabled by the operating system 12a. Wake up could also be more permanently enabled (or disabled) by hardwiring the PME\_En field in the PMCS Register 34. The MAC subsystem 30 uses this field to determine whether wake up is required.

The MAC subsystem 30 transitions the physical layer interface 18 into one of three power states depending upon the power management state of the Gigabit Ethernet device 16 (i.e., D0 Active, D0 Uninitialized, D Reset or D3) and whether wake up has been enabled by the operating system 12a. The three power states of the physical layer interface 18 are defined in Table IV below.

Table IV

| <b>Power State of Physical Layer Interface 18</b> | <b>Definition</b>   |
|---|---|
| High  | Physical Layer Interface 18 is configured to transmit/receive data at full 1,000 Mbps rate. |
| Low   | Physical Layer Interface 18 is configured to transmit/receive data at 100 or 10 Mbps.       |
| Off   | Physical Layer Interface 18 is powered off; not transmitting/receiving any data.            |

When the MAC subsystem 30 transitions the physical layer interface 18 from its high power state to its low power state, the physical layer interface 18 will re-negotiate its data link speed from 1000 megabits per second to 100 or 10 megabits per



second. Similarly, when the MAC subsystem 30 transitions the physical layer interface 18 from its high power state to its off power state, the physical layer interface 18 ceases any transmission or receipt of data.

- 5 Table V illustrates how the MAC subsystem 30 controls power to the physical layer interface 18 in each of the four power management states when wake up is enabled and disabled.

Table V

| <b>Power Management State<br/>of Device 16</b> | <b>Wake Up</b> | <b>Power State of<br/>Physical Layer<br/>Interface 18</b> |
|--|----------------|---|
| D0 Active                                      | Enabled        | High  |
| D0 Active                                      | Disabled       | High  |
| D0 Uninitialized                               | Enabled        | Low   |
| D0 Uninitialized                               | Disabled       | Off   |
| D Reset  | Enabled        | Low   |
| D Reset  | Disabled       | Off   |
| D3   | Enabled        | Low   |
| D3   | Disabled       | Off   |

10 Many computer systems provide for an auxiliary power supply in the event of a loss or power down of a system's main power supply. In a computer system built in accordance with certain specifications, e.g., PCI Bus Power Management Interface  
 15 Specification, Revision 1.1., published December 18, 1998, devices that use the auxiliary power supply are limited to draw

a certain amount of current, e.g., 375 mA, from the PCI bus. A device that draws more than the threshold amount of current may cause a reduction in voltage which could result in incorrect operation of the host computer system or may damage the auxiliary power supply.

The Gigabit Ethernet device 16 is capable of transmitting up to 1000 megabits of data per second, which is a 10-fold increase over 100 megabit Fast Ethernet and a 100-fold increase over 10 megabit Ethernet. While the Gigabit Ethernet device 16 is capable of high-speed data transfer, the physical layer interface 18 of Gigabit Ethernet device 16 can consume significant amounts of power when it is in its high power state (i.e., transmitting/receiving data at 1000 megabits per second). In the event of the host computer system 12 losing or powering down its main power supply 12b, the Gigabit Ethernet device 16 with a fully-powered physical layer interface 18 could draw more than the threshold amount of current, e.g. 375 mA, from the auxiliary power supply 12c and risk incorrect operation of the computer system 12 or damage to the auxiliary power supply 12c. However, when the physical layer interface 18 is operating in its low power state (i.e., transmitting/receiving data at 10 or 100 megabits per second) or its off power state, the physical layer interface 18 will not draw more than the threshold of amount current of the auxiliary power supply 12c.

By automatically changing the power management state of the Gigabit Ethernet device 16 out of D0 (52) and powering down the physical layer interface 18 to a low or off state shortly before the main power supply 12b powers down, the risk of the physical layer interface 18 drawing excessive current from the auxiliary power supply 12c is reduced.

A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention.

For example, the automatic power down technique described may be incorporated in any communications device which has a physical layer interface that may draw current in excess of a threshold amount of current from an auxiliary power supply, such as a wireless Local Area Network (LAN) device.

Additionally, the technique may be implemented in any peripheral device in a communications arrangement that generates a signal indicating that the computer system is powering down its power supply or transitioning to an auxiliary power supply. Communications arrangements that utilize a standard PCI or PCI-X bus and otherwise comply with PCI Bus Power Management Specification 1.1., December 18, 1998, PCI Local Bus Specification, Revision 2.2, December 18, 1998, or PCI-X Addendum to the PCI Local Bus Specification, Revision 1.0A, July 24,

2000, are examples of communication arrangements in which a peripheral device may utilize this technique.

The power states of the physical layer interface may be defined in a number of ways so long as the physical layer interface transitions to a power state in which it operates below auxiliary power source threshold (e.g., 375 mA) when the device detects a signal on the bus indicating that the main power supply is about to go down. For example, the device may be designed to automatically power off the physical layer interface whenever a PCI reset signal is asserted regardless whether wake up has been enabled.

Accordingly, other embodiments are within the scope of the following claims.